Data-driven Optimization and Model Discovery in Markov Processes

Shie Mannor (Technion)

April 19, 2012

Many real-world dynamic decision problems are characterized by uncertainty: the model parameters or even the model representation are not known and have to be estimated from data. In this talk we address modeling and optimization in multi-stage decision problems focusing on Markov decision processes.

We start from the problem of how to use a sample of trajectories to choose from a candidate set of possible state spaces in Markov processes. Standard approaches to solving this problem for static models use penalized maximum likelihood criteria that take the likelihood of the trajectory into account. Surprisingly, these criteria do not work even for simple fully observable finite Markov processes. We propose an alternative criterion and show that it is consistent.

Given that we know how to choose the right model we address the problem of policy optimization in Markov decision processes with parameter uncertainty. We describe several approaches to handling model uncertainty and analyze their complexity and discuss their implications.